SAND HILL
ECONOMETRICS

Measuring Risk and Performance for Private Equity

Susan E. Woodward

Sand Hill Econometrics

SWoodward@SandHillEcon.com

Draft of August 11, 2004

Abstract

Institutional investors find alternative assets such as venture capital and buyout funds appealing for their high returns given the risk they believe they assume. The beliefs about the risk of these portfolios seem to result from the low correlation between reported fund quarterly returns with quarterly returns on stock market indices. The fund quarterly returns, reported by general partners, come from a mix of current and stale company valuations, and thus are not accurate measures of the true change in value of the portfolio for the quarter. I offer an approach to measuring risk and performance which compensates for the stale valuations and shows that the risk of the two major sectors of private equity, venture capital and buyouts, is considerably higher than is generally believed, both in terms of standard deviation of return and *beta*. I find a *beta* for venture capital of 2.0, and for buyouts, .86. Both sectors have correlations with broad equity market indices well above 0.5. With higher risk, the risk-adjusted returns are correspondingly lower.

Introduction[1]

Private Equity is valued as an asset class by institutional investors both for its potential above-average performance and for its power to diversify a portfolio beyond what is available in traded securities. Many institutional investors believe that both venture capital and other private equity funds, held as limited partnerships, have returns that exhibit low correlations with returns on traded equities[2]. Hence, the private equity returns appear high given that much of the risk appears diversifiable. This belief could easily follow from a standard risk-adjusted return analysis of the quarterly returns reported to investors by general partners, and on the averages of these returns as reported by organizations such as Venture Economics[3] and Cambridge Associates. The simple correlation between these quarterly returns and returns on broad market indices such as the S&P500 or Wilshire5000 are low enough to suggest that both of these types of private equity are less risky than plain vanilla stock market exposure. With average returns well above those on the stock market, both types of private equity appear to be superior performers.

I will show that the low estimates of risk and correlation result from a failure to recognize that the values reported by GPs are a mix of current and stale values and to analyze them accordingly. I propose an approach to measuring risk

---

[1] This edition reflects the benefit of comments from John Hand and Morten Sorenson. More comments are welcome.

[2] For example, see the UC Regents web site,

http://www.universityofcalifornia.edu/regents/regmeet/mar03/603.pdf

[3] VE reports a beta for venture capital of .84 in its *2001 Investment Benchmarks Report,* and Ibbonson Associates uses the estimates from Peng Chen and co-authors to obtain a correlation with the public markets of .04.

and performance with the returns reported by GPs to deal with the stale values for some assets in the portfolio.

Earlier approaches to measuring risk and performance for private equity are diverse. Among those who have dealt with fund-level quarterly return data are Chen, Baierle, and Kaplan (2002), Emery (2003), and Gompers and Lerner (2000). Chen and co-authors regress the average quarterly returns reported by Thomson Venture Economics on market returns and report a *beta* of .0034 and a correlation with the public markets of only .04. Emery comes the closest to the approach used here. He sees that stale pricing would result in serially correlated fund returns. He uses regressions with lagged values of a benchmark to establish the presence of serial correlation in fund returns. His preferred solution to measuring risk is to use longer time intervals—years rather than quarters—for measuring return. His estimates private equity risk and correlation with the stock market are considerably lower than the ones presented here.

Gompers and Lerner propose a method for evaluating particular partnerships. They address the problem of stale pricing in the computation of quarterly returns by bringing each company's last deal price forward in time using the industry-specific sub-index of Nasdaq, re-calculate the quarterly returns using updated company prices, and then regress these "refreshed" returns on market returns to calculate an *alpha* and a *beta* for the partnership. The usefulness of this approach depends on how well the Nasdaq sub-index matches the companies in the portfolio.

Several other approaches use data on comparable public companies. To measure risk for venture capital, Kerins, Kiholm Smith, and Smith (2003) construct a portfolio of newly public companies by matching venture companies as closely as possible to companies newly public. For this portfolio they obtain a *beta* of about one. The clear question that arises for such an approach is whether newly public companies are indeed good risk matches for private companies. The results here suggest that private companies are riskier than public ones. Lundqvist and Richardson (2003) take a "public market equivalent" approach that was introduced by Long and Nickles (1995). They use cash flow data from funds representing a substantial fraction of US private equity, and measure both risk and performance for a mixed portfolio of venture capital funds and buyout funds, also by matching portfolio companies to a Nasdaq industry index. They then treat the cash flows into and out of these funds as investments and sales in the relevant index, carry interim balances forward at riskless rates, and treat the resulting ending value as a measure of performance. They find excellent performance (positive *alpha*) for private equity. Again, the question arises, how good are the indices for public companies at matching the risk of those that are private or taken private? If we guess too low on the *beta*, the estimate of *alpha* will be too high. Kaserer and Diller (2004) use the public market equivalent approach on a large portfolio of European private equity (buyout type) funds.

Cochrane (2002), Peng(2001), and Quigley and Woodward (2002) all use company-level deal data (private rounds of funding, IPOs, acquisitions, and shut downs) from Sand Hill Econometrics. Cochrane takes a maximum likelihood approach to an early subset of the SHE data and gets a *beta* of about one and finds substantial risk-adjusted returns. Peng uses a method-of-moments repeat-sales approach with a re-weighting procedure to correct for selection bias. Quigley and Woodward use a repeat-sales approach, corrected for the selection bias present in the sharing of post-money value, and construct an index whose *beta* is then measured against the Wilshire 5000. Sand Hill Econometrics itself, in building the Sand Hill Index, does not use classic repeat-sales and also does not use Heckman-type bias correction, mainly because the observed distributions of return depart so far from log-normal. Instead, Sand Hill estimates value for every date (month) for every company that is alive (between its first round and its exit) using observed valuations in that month and on either side of it, and corrects for bias by comparing censored (biased) data with uncensored (unbiased) data. For more details on the building of the index, see Woodward and Hall (2003).

The Stale Pricing Problem

In the stock market, some stocks trade more frequently than others. For the most thickly traded stocks, we can be fairly certain that the closing price is for a trade that took place at 4 pm Eastern Time, the closing time for the market. For thinly traded stocks, the last price may be from a trade earlier in the day, from a trade the previous day, or even some days in the past. When this is the case, regressing returns for a stock on returns on a market index over identical time intervals will tend to underestimate the *beta* on the stock because the intervals over which return is measured are not synchronous. If our estimate of *beta* is too low, then our estimate of *alpha* will be too high; this can be proven as a statistical proposition, but the intuition in this application is that if we attribute too little return to risk, then too much is attributed to pure performance.

The problem occurs for both thinly traded stocks, which appear to lag the market, and also for the most thickly traded stocks, which appear to lead it. A straightforward method for addressing this problem was suggested by Dimson (1979) and Scholes and Williams (1977). The adjustment needed is to include leading and lagging values of market returns as well as the contemporaneous return, and sum the resulting coefficients to get the correct *beta.* For applications involving estimating the *beta* for individual traded stocks, two days of leading returns and two days of lagging returns are usually sufficient to capture all of the true correlation between one stock's returns and overall market returns. Because the daily returns on the broad index are uncorrelated with (orthogonal to) each other, it is possible to determine the correct number of leads and lags simply by including enough so that the farthest out are essentially zero (coefficients and t-statistics close to zero), then dropping the leads and lags with coefficients close to zero. I find that the lags are well-behaved and decline in an orderly fashion.

As an illustration of the problem, consider the relationship between the S&P500 and the Wilshire5000, two much-used indices of the US stock market. The S&P500 tracks the prices for 500 companies that are strategically selected by Standard & Poor's to represent the stock market. The Wilshire5000 tracks all US publicly traded stocks. Many stocks in the Wilshire5000 trade infrequently. This is evident even using monthly returns in regressions comparing the Wilshire5000 and the S&P500. Consider the results of two regressions. Table 1 shows the relation between the Wilshire5000 and the S&P500 looking only at the contemporaneous returns, and the second includes a leading and a lagging value of the S&P. From the first regression, it appears that the indices are similar but that the average return on the Wilshire is slightly higher (positive *alpha*). I use one-month Treasury Bills from Ibbotson for the riskless rate.

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.00133	*0.00082*	1.62
beta	**0.99494**	*0.01818*	54.73
R^2	0.96		
Standard Deviation, monthly Wilshire 5000 return premium	.04546		

Table 1. Regression of the montly return premium for the Wilshire 5000 on the return premium for the S&P 500, contemporaneous returns only, Jan, 1994 to May, 2004.

When we add one leading and one lagging value of monthly return premiums on the S&P500 as independent variables, we get an enriched view of the relationship between the two indices:

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.001	*0.001*	1.33
SP500 return premium, one month ahead	0.034	*0.018*	1.90
SP500 return premium	0.995	*0.018*	55.3
SP500 return premium, one month lagged	0.025	*0.018*	1.37
Sum, *beta*	**1.054**	*.0318*	33.09
R^2	0.96		
Standard Deviation, monthly Wilshire return premium	.0456		

Table 2. Regression of the monthly return premium for the Wilshire 5000 on the return premium for the S&P 500, contemporaneous return plus one month leading and one month lagging, Jan. 1994 to May 2004

Adding leading and lagging returns to the regression raises the estimate of *beta* for the Wilshire on the S&P from .99 to 1.05. The *t*-statistic[4] for summed coefficients on the leading and lagging variables indicates that it is very unlikely that we would obtain coefficients this high if no such leads and lags were truly present. Evidently "the" market is mainly comprised of stocks that are traded not quite absolutely continuously. Thus, the stocks that *are* traded continuously cause the leading value on the Wilshire to explain a bit of variation in the S&P.

The coefficient on the leading value of the S&P indicates stale pricing in the S&P, while the coefficient on the lagged value of the S&P indicates stale pricing in the Wilshire. So which index is "the" market? Neither. On the one

[4] We use a *t*-statistic on the summed coefficients rather than a marginal *F* because it is a more powerful test for whether the addition to the *beta* is different from zero.

hand, the Wilshire5000 includes many stocks that the S&P500 does not, and captures detectable additional variation as a result. But some of this variation is coming from stale pricing. You cannot buy the Wilshire5000 at the value of the Wilshire5000, because some of the prices used to calculate it no longer prevail. As analyses of individual stocks show, the S&P500 also has some stale prices,[5] so the S&P500 cannot be bought at its reported value either, except as an exchange-traded fund or a futures contract, which is not quite the same thing. Compared to the Wilshire, the S&P has less stale pricing, but the Wilshire captures variation in the value of the smalls stocks in a way the S&P does not. No one has attempted to construct an index of US stocks that represents all of the stocks in the market at their *current* prices using estimates of current price where current prices are unavailable.

The point of this exercise is to demonstrate that detectable non-synchronous trading effects are present even in monthly returns in markets that are very familiar, where the larger stocks trade every moment. Acknowledging that the Wilshire5000 is not a perfect index of "the market", we will nonetheless use the Wilshire because it is broader and we want our investigation of private equity to capture as much about the variation in value of small public companies as possible.

The problem of measuring risk and performance in the presence of non-synchronous trading arises in other contexts where the problem is much more acute. This was recognized by Asness, Krail, and Liew (2001) in assessing the performance of hedge funds. Hedge funds hold many securities that are seldom traded. In addition, because hedge funds are organized so as to not be subject to

[5] This technique was applied to individual company estimates of *beta* for small, thinly traded stocks by Ibbotson, Kaplan, and Peterson (1997), who found that the simpler regressions under-estimate individual company *betas*.

1940 Investment Act requirements (and are thus sold only to "qualified" investors), hedge fund managers have more freedom in reporting values and returns to investors. Asness and co-authors show that including not only contemporaneous monthly returns, but also returns for three lagging months, doubles the hedge fund *betas*, raises the correlation of their returns with the market from .5 to .7, and changes the measure of average risk-adjusted excess return, *alpha*, from a positive to a negative value for hedge funds overall. Getmansky, Lo, and Makarov (2003) also find similar staleness in hedge fund pricing.

Application to Venture Capital

Venture capital partnerships make investments in young private companies. In their earliest rounds of funding, these companies are generally without revenues. The hope and plan for these companies is to become or be acquired by a public company.

The investors generally give the companies only enough funding to see them through the next year or two, not more. If the results from this investment are such that the company is still promising, a new round of investment is made. Typically developing companies do four, five, or more rounds of funding between the first round in which they solicit money from outside investors (beyond the founder, family and friends who invest in their seed round) and their exit via IPO, acquisition, or shut down.

A new security is issued for each round of investment. Usually the founders, and possibly the seed round investors, hold common stock. The stock issued at the first formal round to outside investors is usually referred to as the A Round (or A Series) or First Round Convertible Preferred stock. The *preference* enjoyed by the preferred shareholders is typically that in the event of a liquidation that is not sufficiently attractive for all investors to choose to convert their preferred shares to common shares, the holders of the preferred stock recover their entire investment (or some multiple of it) before the common shareholders receive any value. This preference mainly serves to prevent the managers of the firm from paying a liquidating dividend to all shareholders with the outside investors' money, leaving the outside investors with only a fraction of what they invested. The structure of each subsequent investment is similar—the most recent investors have a liquidation preference over prior investors.

Many investors and venture analysts believe that the liquidation preferences make the set of equity-like securities issued to finance the companies a set of nested call options on the assets of the company. In our experience, this is only a vague approximation of the actual rights conferred by the preferences. The preferences are potentially of importance only in a small fraction of the possible outcomes of the investments in startup companies. In about half of the outcomes of such investments, the company expires worthless and no investors convert, making the liquidation preferences irrelevant. In about 30 percent of the outcomes, there is a public offering or acquisition sufficiently attractive that all investors convert to common, and again, the preferences are irrelevant. In the cases of disappointing acquisitions, perhaps as much as 20 percent of the outcomes, the preferences can potentially influence final payoffs.

In these situations, it is not unusual for the most senior shareholders to need cooperation from the most junior, typically the common shareholders, to accomplish the disappointing acquisition. To get the required total of votes, they may need the cooperation of the less senior preferred classes as well. As a result of this dynamic, essentially all of the preferences are up for negotiation in the state of the world in which the preferences can potentially have an impact on the distribution of value across shareholders. We conclude that preferences clearly matter in protecting shareholders from a confiscatory liquidation at the time of their investment, and they are useful for tax purposes—see Gilson and Schizer (2003)—but their ultimate impact on investment distributions in mediocre outcomes are potentially not closely aligned with the preferences and in any case not well studied. The likely irrelevance of preferences for ultimate investment outcomes is important in interpreting of the valuation methods used by GPs, who also ignore preferences.

Each quarter general partners report the value of each company in the fund to the limited partners. These values are nearly always based on each company's most recent round of funding. The GPs usually take all of their stock holdings in a company, common shares, plus A preferred shares, plus B preferred shares, and multiply all by the price of the stock in the most recent round (for example the B round here), ignoring any potential option value to the preferences. This price may have been determined at a round of funding that is just completed, and is thus a recent price, but also it may be some months or years old.

For functioning companies, departures from the practice of valuing each company at its most-recent-round price are rare. For companies for whom failure appears likely, we sometimes see GPs who still carry the company at its most recent round price, some who write it down partially, and some who write it down completely even though the web site is still up and the company still has employees. So as a general matter, the values that GPs report to LPs each quarter are a mix of current and stale valuations, and there is mainly consistency in the reporting of value except for failing companies.

The presence of staleness in GP valuations implies that if there is any relation between stock market activity and the prices found for investments in private companies, the quarterly changes in value reported by GPs will be related not only to current returns on the stock market, but also to past returns. A regression of portfolio quarterly returns on index quarterly returns will understate the relation, as such a regression does for thinly traded stocks and hedge funds. The appropriate correction is essentially the same as that employed by Dimson to measure *betas* of public companies and by Asness and co-authors to measure hedge fund *betas*: include (leading and) lagging values of market returns in the regression, and add up the coefficients. It is Dimson's adjustment on a grand scale, intended as a solution to a grander version of the same problem, non-synchronous trading.

The phenomenon is clear in the comparison of two regressions of venture capital returns on stock market returns using quarterly returns from venture funds reported by Cambridge Associates and the Wilshire5000 for the stock market. The venture returns from CA are averages of returns from a set of funds that represent a substantial fraction of all venture capital, but not the universe. This first regression includes only the contemporaneous return on the stock market as an independent variable. Because we are using quarterly data, the three-month Treasury bill is used as the riskless rate of interest.

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.0195	*0.014*	1.40
beta	**0.593**	*0.15*	3.84
R^2	0.17		
Standard Deviation, quarterly CA return premium	0.128		

Table 3. Regression of the quarterly return premium for the Cambridge Associates venture fund return series on the return premium for the Wilshire 5000, contemporaneous returns only, 1985-2003. Three-month Treasury Bills are used for the riskless rate.

These results make venture capital look like a money printing machine. The *beta* is well below one, and the *alpha* is nearly 2 percent per quarter. The standard deviation of annualized venture returns is roughly double that of the stock market, but the regression results suggest that this risk is nearly all diversifiable.

The Durbin-Watson statistic on this regression should give us pause, however, because it suggests the presence of serial correlation in the error terms. This will arise if there is serial correlation in the quarterly venture returns but not in the quarterly stock index returns because the stock market returns are not serially correlated and hence cannot explain the serial correlation. We generally expect that in an efficient market, there will be very little serial correlation in returns. The following regression, which includes a correction for autoregressive error terms, shows a clear pattern of serial correlation in the CA venture returns:

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.039	*0.034*	1.136
CA return, lagged 1 quarter	-0.073	*0.394*	-0.185
Lagged 2 quarters	0.251	*0.241*	1.041
Lagged 3 quarters	0.296	*0.116*	2.551
Lagged 4 quarters	-0.188	*0.129*	-1.460
Lagged 5 quarters	-0.201	*0.114*	-1.759
Autoregressive parameter	0.540	*0.398*	1.356
R^2	0.45		
Standard Deviation, CA return premium	0.1218		

Table 4. Regression of the quarterly return premium for the Cambridge Associates venture fund return series on its own past returns, using five lagged quarters, 1985-2003

Given this degree of serial correlation, the lagged market returns are likely to be important in revealing the full story of risk and performance. Contrast the regression above in Table 3, using only the contemporaneous stock market return, to the one below in Table 5, where lagging values of stock market quarterly returns are included as independent variables:

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	-0.003	*0.015*	-0.175
Wilshire return premium	0.684	*0.152*	4.496
Lagged 1 quarter	0.276	*0.153*	1.806
Lagged 2 quarters	0.340	*0.150*	2.264
Lagged 3 quarters	0.375	*0.153*	2.453
Lagged 4 quarters	0.302	*0.152*	1.982
Lagged 5 quarters	0.082	*0.152*	0.538
sum, *beta*	**2.059**	*0.566*	3.64
R^2	0.45		
Standard Deviation, CA quarterly return premium	0.125		

Table 5. Regression of the quarterly return premium for the Cambridge Associates venture fund return series on the return premium for the Wilshire 5000, contemporaneous return plus five lagging quarters, 1990-2003

By this assessment, venture capital looks less attractive. The *beta* is 2.06, and the *alpha* is now *minus* 0.3 percent per quarter. The t-statistic suggests that it is very unlikely that the lagged returns on the market are unrelated to the CA venture returns.

These results should not be taken as the last word on venture performance. First, the quarterly returns from Cambridge Associates' averages of partnership returns are not value-weighted. If small funds perform on average worse than large ones, regressing equal-weighted (instead of value weighted) averages on a market index can make all of venture capital look like a losing investment. Second, these results come only from funds who have voluntarily shared data with Cambridge. We might suspect it is the better performers who are likely to share results. The CA fund returns are net of fund expenses and carried interest.

We therefore turn to an alternative benchmark for venture capital, the Sand Hill Index. The Sand Hill Index is constructed to be the venture analog of a comprehensive public company index. This index is built from company valuation events (private rounds of funding, IPOs, acquisitions, and shut downs), not from fund returns, by a repeat-sales method. Note that the Sand Hill Index is not constructed from returns on funds, but from returns on companies held by funds. The companies included in the index are close to the universe of private companies that have sought funding from outside investors. While nearly all companies make an effort to publicize their fund-raising events, and thus are captured in the data, only some, generally the more successful, share the post-money value implied by the terms of a fund raise. Less successful companies also often do not share with the public a valuation for a disappointing acquisition. The Sand Hill Index estimates value for companies that do not share value publicly with methods that correct for the upward bias in reporting. A more complete discussion of the construction of this index is in Woodward and Hall (2003).

We build the Sand Hill Index so as to give the current value, each month, for a continuously reinvested portfolio all venture companies in proportion to their market values. Because we build the index from observations of value that are months and sometimes years apart, and therefore interpolate value for intervening months, there is considerable serial correlation in the returns on the Sand Hill Index. We thus also include an autoregressive correction factor. Apart from the serial correlation induced by interpolation, another force causing serial correlation is that many venture-type companies do a round of funding in multiple closings over many months, all at the same price. In addition, the negotiation of price for a round usually takes time measured in months. The high cost of finding price may introduce inertia into the time series of prices. We thus include lagging values of market returns to estimate *beta* even when using the Sand Hill Index. We examine all returns as excess returns over 30-day Treasury Bills. The Sand

Hill Index is compiled on a monthly basis, so we use monthly returns in the next regressions of the Sand Hill Index excess returns on Wilshire5000 excess returns. Like equity market indices, the Sand Hill Index is compiled gross of fund-level expenses and carried interest.

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.00031	*0.00215*	0.14263
Wilshire return premium	0.90709	*0.04511*	20.11032
Lagged 1 month	0.02844	*0.04521*	0.62900
Lagged 2 month	0.05362	*0.04484*	1.19581
Lagged 3 months	0.07212	*0.04480*	1.60960
Lagged 4 months	0.04542	*0.04461*	1.01812
Lagged 5 months	0.08649	*0.04459*	1.93967
Lagged 6 months	0.03888	*0.04486*	0.86671
Lagged 7 months	0.01319	*0.04490*	0.29374
Lagged 8 months	0.01885	*0.04493*	0.41952
Lagged 9 months	0.04150	*0.04457*	0.93097
Lagged 10 months	0.03047	*0.04469*	0.68185
Lagged 11 months	0.06748	*0.04478*	1.50697
Lagged 12 months	0.06396	*0.04490*	1.42432
Lagged 13 months	0.10819	*0.04492*	2.40837
Lagged 14 months	0.06515	*0.04449*	1.46443
Lagged 15 months	0.09800	*0.04172*	2.34868
Lagged 16 months	0.11002	*0.04195*	2.62286
Lagged 17 months	0.06464	*0.04212*	1.53466
Lagged 18 months	0.08653	*0.04225*	2.04810
sum, *beta*	**2.000013**	*0.18450*	10.84
R^2	.74		
Annual Standard Deviation, Sand Hill return premium	0.489		

Table 6. Regression of the monthly return premium for the Sand Hill Index on the monthly return premium for the Wilshire 5000, contemporaneous return plus eighteen lagging months, 1990-2003.

This regression, which uses returns that are constructed to represent changes in value for all of venture capital on a current value, continuously-invested, value-weighted basis, gives a similar picture of venture capital in terms of risk with a total *beta* of 2.0, statistically not different from the 2.05.[6] In terms of return, however, the difference is more profound than it first appears because the CA returns are *net* of expenses and carried interest, while the Sand Hill returns are *gross*. The Sand Hill *alpha* of 3 basis points per month amounts to about .36 percent per year. A typical structure for GP compensation in a venture partnership is 2 percent per year in expenses and 20 percent carried interest (20 percent of any returns over zero). Amortized over 10 years, the carried interest is roughly 2 percent per year, for a total of 4 percent, which would leave the risk-adjusted performance of venture after expenses and carry decidedly negative for this period (1989-2004). It appears that the performance of venture capital in general is not so outstanding as was suggested by the first simple regressions of contemporaneous returns. It also appears that the reward to venture investing is not quite sufficient to cover the systematic risk, much less leave a further premium for illiquidity.

The Sand Hill Index leads the Cambridge returns out a year (four lagged values of returns on the Sand Hill Index bear a non-random relation to CA return premiums). This results not because the GP returns are inherently predictable, but because they are comprised partly of stale valuations. Statistical elimination of most of the staleness thus produces a series that predicts the part stale series.

There are reasons to suspect that the lag structure of returns for venture capital is not entirely stable. When the market was in its upswing in the late 1990s, the time elapsed between rounds of funding for venture companies was

---

[6] Using quarterly returns on the Sand Hill Index and the Wilshire5000, we get a *beta* of 2.03.

several months shorter than the time elapsed during the period 2000-2002. When values are rising, the general partners have an incentive to do a new round of funding so that they can "print a trade" at a higher price. When the market is falling, they might delay a round of funding to delay printing a trade at a lower value. In addition, the delay in writing off failed companies may be influenced by whether there is some good news to offset the bad.

Buyout Funds

Buyout funds are, like venture capital, considered as part of the general class of private equity. Both are typically organized as limited partnerships with a general partner who provides some of the funding for the partnership and makes the investment decisions for the fund. But there are important differences. In venture funds, the companies are usually startups with no revenues. In buyout funds, the companies are usually mature public or private companies that have been sustained by their own revenues for many years. Venture investments are most often syndicated—several venture funds will invest in each round of fund raising. In buyouts, we more often see one fund make the entire investment in a single company. Venture funds typically advise portfolio companies and sit on its board of directors, while buyout funds are more likely to be thoroughly involved in day-to-day management. Venture companies are not usually given enough funding to continue operating for more than a year or so, and if at the end of that period, the venture still seems promising, the company does another round of fund raising. It is not unusual to see four, five or more rounds of funding from startup to exit. Each fund raise is a pricing event, whose price is used to report value by venture GPs. Buyout funds, by contrast, typically buy a company, sustain it on its own revenues while making changes in its activities or organization, and sell it two to seven years later. Quarterly valuations reported by GPs in the interim are their own internal valuations, not values from actual transactions. About half of all venture-funded companies fail. Few buyout companies do. The table below summarizes these differences:

	Venture	Buyouts
Portfolio Company stage	*early*	*mature*
Company revenues	*none*	*sustaining*
Number of fund investors	*many*	*one*
Management involvement	*moderate*	*heavy*
Number of cash infusions by investors	*4-7*	*1*
Use of debt by investor	*never*	*nearly always*
Source for quarterly valuations	*funding round*	*internal*
Ultimate company failure rate	*about 50%*	*rare*

An important implication of the absence of syndication and new pricing events between initial investment and realization for buyouts is that the general partner has considerable freedom in reporting the value of the company to investors between the date of the original investment and the date when the company is sold. This freedom extends to having the option to sell a portfolio company and book a realized value. Of course, the last reported return must reconcile the earlier estimates of value with the final realized value. Thus, while the precise lag structure for buyout valuations is somewhat more difficult to predict than the structure for venture capital, we should not be surprised if including lagged returns increases the estimate of risk for buyout fund investments.

The buyout and venture results are indeed similar in that the lagged returns are important for measuring risk, but buyouts funds are considerably less risky than venture funds and staleness appears to be less severe than for venture funds.

The next two tables report the results for quarterly returns on US buyout funds from Thomson Venture Economics, 1988-2004, Q1. These returns represent 38 percent of US buyout funds, and about 68 percent of the dollars in US buyout funds. The first table shows the results from regressing the buyout fund returns (return for all funds for the quarter) on the contemporaneous Wilshire5000, and the next the results from including lagging quarters of the index.

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.012	*0.006*	2.099
Wilshire return premium	**0.389**	*0.070*	5.589
R^2	0.34		
Standard Deviation, buyout fund return premium	0.055		

Table 7. Regression of the quarterly return premium for the VE buyout return series on the return premium for the Wilshire 5000, contemporaneous returns only, 1990-2003

A second estimation, reported below, shows the difference from including lagging returns on the index.

Variable	*Coefficient*	*Standard Error*	*t-Statistic*
intercept, *alpha*	0.005	*0.006*	0.805
Wilshire return premium, lagged 1 quarter	0.384	*0.066*	5.826
Lagged 2 quarters	0.135	*0.067*	2.026
Lagged 3 quarters	0.127	*0.067*	1.903
Lagged 4 quarters	0.113	*0.064*	1.777
Lagged 5 quarters	0.103	*0.063*	1.633
sum, *beta*	**0.862**	*.2368*	3.64
R^2	0.465		
Standard Deviation, buyout return premium	0.0553		

Table 7. Regression of the quarterly return premium for a buyout portfolio on the return for the Wilshire 5000, contemporaneous returns plus six lagging quarters, 1990-2003

This second equation shows that buyout funds are, risk/return-wise, similar to general stock market exposure, with a *beta* of .86, a correlation overall with the market of about .7 (the square root of the adjusted R^2 on the regression equation), and a standard deviation of return of roughly 23 percent on an annualized basis. Again here, the quarterly returns have so much serial correlation that the annual standard deviation of return cannot be found just by doubling the quarterly standard deviation. While the *alpha* is positive, it is smaller than its own standard error. And since these returns are voluntarily shared with Thomson VE, and not a random sample or the universe, they may be biased upwards.

Investing in buyout funds may add value for investors because of the manager's skill in choosing and managing portfolio companies, but the nature of the risk is not fundamentally different from investing in the stock market. A wise and cynical person once wryly claimed that buyout funds were just S&P500 exposure, liberated from marking to market. In light of these results, it is hard to disagree.

The underlying behavior that produces the lag structure for buyout returns is not as straightforward to identify as that for venture, where actual transaction prices are used, however stale. Because so much time passes from the buyout to the sale of a company in a buyout fund, the GPs usually estimate its value for quarterly reporting. They could be using stale comparables, or they could be smoothing earnings. Both should produce serial correlation in returns and thus cause the returns to be related to lagging as well as contemporaneous returns on public markets.

Concluding Remarks

Investors receive too optimistic an impression from the standard analysis of performance of buyout and venture funds. Reported correlations with the stock market are biased downward and risk-adjusted returns (*alphas*) are biased upward. Bias arises because the period over which the buyout or venture return is measured is not the same as the period over which the stock market returns are measured. The lack of synchronicity results in low estimated correlations. Investors get the false impression that buyout and venture investments are not very risky and that they out-perform other investments on a risk-adjusted basis.

Up-to-date valuations are difficult to come by, so the valuations actually reported by the general partners of venture and buyout funds are often either values observed at some point in the past, in the case of venture capital, or are estimates of value, in the case of buyouts. The use of these values in a standard return regression is dangerously misleading. This paper offers a simple and practical approach to measuring risk and assessing performance that corrects the bias arising from reporting practices.

Our approach is to include enough lagged returns on the index or benchmark as independent variables in the standard performance regression to capture all of the correlation between a broad stock market index and the reported returns on the asset, and then summing the coefficients to get the risk measure. In the case of venture capital, the *beta* measure of risk increases dramatically, from .6 to 2.0. The estimated *alpha* is a misleading 1.8 percent of excess return per quarter in the uncorrected regression and essentially zero in the corrected one. For buyouts, *beta* is a low 0.4 with an *alpha* of 1.2 percent per quarter in the biased regression and rises to 0.9 with an *alpha* of 0.5 percent per quarter in the corrected regression.

References

Asness, Clifford, Robert Krail, and John Liew, (2001) "Do Hedge Funds Hedge?" *Journal of Portfolio Management*, Fall

Chen, Peng, Gary T. Baierl, Paul D. Kaplan, (2002) "Venture Capital and its Role in Strategic Asset Allocation" *Journal of Portfolio Management* Winter

Dimson, Elroy, (1979) "Risk Measurement when Shares are Subject to Infrequent Trading" *Journal of Financial Economics*, 7

Emery, Kenneth, (2003) "Private Equity Risk and Reward: Assessing the Stale Pricing Problem" *Journal of Private Equity*, Spring

Gilson, Ronald, and David Schizer, (2003) "Understanding Venture Capital Structure: A Tax Explanation for Convertible Preferred Stock" *Harvard Law Review*, January.

Getmansky, Mila, Andrew Lo, and Igor Makarov, (2003) "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns" *MIT Working Paper 4288-03*, March.

Kerins, Frank, Janet Kiholm Smith, and Richard Smith, (2003) "The Opportunity Cost of Capital for Venture Capital Investors and Entrepreneurs" February, working paper.

Ibbotson, Roger, Paul Kaplan, and James Peterson, (1997) "Estimates of Small Stock Betas are Much Too Low" *The Journal of Portfolio Management*, 24, pages 104-111.

Long, Austin M., and Craig Nickels, (1995) "A Method for Comparing Private Market Internal Rates of Return to Public Market Index Returns" manuscript, University of Texas System, August 28.

Lundqvist, Alexander, and Matthew Richardson, (2003) "The Cash Flow, Return, and Risk Characteristics of Private Equity" *NBER Working paper #9495*, January.

Peng, Liang, (2001) "Building a Venture Capital Index" *Yale Working Paper 00-51*, August.

Kerins, Frank, Janet Kiholm Smith, and Richard Smith, (2003) "Opportunity Cost of Capital for Venture Capital Investors and Entrepreneurs" working paper, February.

Scholes, M., and J.T. Williams, (1979) "Estimating Betas from Non-Synchronous Data, *Journal of Financial Economics*, 5

Woodward, Susan E., and Robert E. Hall, (2003) "Benchmarking the Returns to Venture" *NBER Working Paper 10202*, December.